AFRL-IF-RS-TR-2002-98
Final Technical Report
May 2002

PROTOCOLS FOR SECURE SURVIVABLE ACTIVE INTERNETWORKING

University of California at Santa Cruz

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. F320

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2002-98 has been reviewed and is approved for publication.

APPROVED: Priscilla Cassidy

PRISCILLA CASSIDY
Project Engineer

FOR THE DIRECTOR:

WARREN H. DEBANY, Technical Advisor
Information Grid Division
Information Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

1. AGENCY USE ONLY (Leave blank)

2. REPORT DATE
MAY 2002

3. REPORT TYPE AND DATES COVERED
Final Jun 97 – Oct 99

4. TITLE AND SUBTITLE
PROTOCOLS FOR SECURE SURVIVABLE ACTIVE INTERNETWORKING

5. FUNDING NUMBERS
C - F30602-97-1-0291
PE - 62301E
PR - F320
TA - 00
WU - 01

6. AUTHOR(S)
J. J. Garcia-Luna-Aceves

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of California at Santa Cruz
Computer Engineering Department
1156 High Street
Santa Cruz California 95064

8. PERFORMING ORGANIZATION REPORT NUMBER
N/A

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington Virginia 22203-1714

AFRL/IFGA
525 Brooks Road
Rome New York 13441-4505

10. SPONSORING / MONITORING AGENCY REPORT NUMBER
AFRL-IF-RS-TR-2002-98

11. SUPPLEMENTARY NOTES
AFRL Project Engineer: Priscilla Cassidy/IFGA/(315) 330-1887/Priscilla.Cassidy@rl.af.mil

12a. DISTRIBUTION / AVAILABILITY STATEMENT
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 Words)*

This report covers work done in three areas:

1) Active destination-oriented QoS support;
2) Active and secure routing and multicasting;
3) Trusted dissemination of active packets.

14. SUBJECT TERMS
Protocols, Active Networks

15. NUMBER OF PAGES
18

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT
UNCLASSIFIED

18. SECURITY CLASSIFICATION OF THIS PAGE
UNCLASSIFIED

19. SECURITY CLASSIFICATION OF ABSTRACT
UNCLASSIFIED

20. LIMITATION OF ABSTRACT
UL

Contents

1. INTRODUCTION

Supporting real-time multimedia applications in such dynamic environments as the joint tactical internet cannot be done simply by applying existing Internet protocols and architectures. First, today's Internet routing and multicasting protocols provide few mechanisms, if any, to protect the exchange of control information or the provision of qualities of service to user applications. By contract, in the tactical internet, nodes and links may be compromised and routing protocols must adapt to drastic changes in link quality and connectivity much more frequently than in the quasi-static routing structure of ATM networks and the IP Internet. Second, all the approaches proposed to date for supporting quality of service in IP or ATM internetworks are based on establishing connections (i.e., associations of sources and destinations for which resources are allocated by nodes in the internetwork) in one way or another [15, 14, ?]; in contrast, the constituency and resources of the path supporting a connection from source to destination or a pre-established multicast routing tree cannot be guaranteed in the tactical internet. We propose to develop new communication protocols for a secure, survivable, and active internetworking architecture in which "active packets" are used to modify the behavior of nodes or inject new services. Active packets can contain data, invocations to procedures, or control procedures. They allow the "state" of programmable nodes to be modified proactively to take advantage of knowledge of the environment and application requirements. The University of California at Santa Cruz (UCSC) addressed these challenges by focusing on the following topics:

- *Active destination-oriented QoS support:* We developed a new Internet protocol architecture to provide end users with different qualities of service, without maintaining connections inside the network.
- *Active and Secure Routing and Multicasting:* We developed new protocols for secure and active routing and unreliable and reliable multicasting.
- *Trusted Dissemination of Active Packets:* We investigated techniques to disseminate reliably active packets that modify the behavior of nodes. An active packet dissemination protocol was started that could be used as the building block for trusted interaction among active (i.e., programmable) nodes.

The research work in this project resulted in 11 refereed papers published in journals and conferences, and three Ph.D. theses.

The theses completed with support from this project are the following:

1. Srinivas Vutukury, "Multipath Routing Mechanisms for Traffic Engineering and quality of Service in The Internet," Ph.D. Thesis, Computer Science, University of California, Santa Cruz, March 2001.
2. Brian Levine, "Supporting Large-Scale Group Communication Applications of The Internet," Ph.D. Thesis, Computer Engineering, University of California, Santa Cruz, June 1999.

3. Clay Shields, "Secure Hierarchical Multicast Routing and Multicast Internet Anonymity," Ph.D. Thesis, Computer Engineering, University of California, Santa Cruz, June 1999.

The published articles describing the results of our research in this project are the following:

1. S.Vutukury and J.J. Garcia-Luna-Aceves, "A Practical Framework for Minimum-Delay Routing in Computer Networks", *Journal of High Speed Networks*, Vol. 8, No. 4, pp. 241-263, Wiley, 1999.
2. S.Vutukury and J.J. Garcia-Luna-Aceves, "A Multipath Framework Architecture for Integrated Services", *Proc. IEEE Globecom 2000*, San Francisco, California, November 27 – December 1, 2000.
3. S.Vutukury and J.J. Garcia-Luna-Aceves, "A Traffic Engineering Approach based on Minimum-delay Routing", *Proc. IEEE IC3N 2000*, Las Vegas, Nevada, October 16-18, 2000.
4. B. Levine, J. Crowcroft, C. Diot, J.J. Garcia-Luna-Aceves, and J. Kurose, "Consideration of Receiver Interest for IP Multicast Deliver", *Proc. Infocom 2000*, Tel-Aviv, Israel, March 26-30, 2000.
5. S.Vutukury and J.J. Garcia-Luna-Aceves, "A Distributed Algorithm for Multipath Computation", *Proc. IEEE Globecom '99*, Rio de Janeiro, Brazil, December 5-9, 1999.
6. S.Vutukury and J.J. Garcia-Luna-Aceves, "A Scalable Architecture for Providing Deterministic Guarantees", *Proc. IEEE IC3N 99*, Boston, Massachusetts, October 11-13, 1999.
7. S.Vutukury and J.J. Garcia-Luna-Aceves, "An Algorithm for Multipath Comutation using Distance-Vectors with Predecessor Information", *Proc. IEEE IC3N 99*, Boston, Massachusetts, October 11-13, 1999.
8. G. Denker, J.J. Garcia-Luna-Aceves, J. Meseguer, P.C. Olvecsky, J. Raju, B. Smith, and C.L. Talbot, "Specification and Analysis of a Reliable Broadcasting Protocol in Maude", *Proc. 37th Allerton Conference on Communications, Control, and Computing*, September 22-24, 1999.
9. C. Shields and J.J. Garcia-Luna-Aceves, "A Scalable Protocol for Secure Multicast Routing", *Proc. ACM SIGCOMM 99*, Cambridge, Massachusetts, September 1-3, 1999.
10. S.Vutukury and J.J. Garcia-Luna-Aceves, "A Simple Approximation to Minimum-Delay Routing," *Proc. ACM SIGCOMM 99,* Cambridge, Massachusetts, September 1-3, 1999.
11. J.J. Garcia-Luna-Aceves, S.Vutukury and W.T. Zaumen, "A Practical Approach to Minimizing Delays in Internet Routing Protocols", *Proc. IEEE ICC '99*, Vancouver, Canada, June 6-10, 1999.

This final report is organized as follows. Section 2 presents our work on routing over multiple paths, minimum-delay routing, and new approaches for providing performance guarantees in a scalable manner in computer networks. Section 3 presents our work on multicast routing architectures and protocols. Section 4 presents our work on the reliable dissemination of packets that can be used for down-loading code or instructions to routers. In each of these sections, we summarize the main results of our work, followed by the main papers describing the technical details of our research. Section 5 summarizes directions for future research.

2. MULTIPATH ROUTING AND QUALITY OF SERVICE

The ability to route packets over multiple paths becomes essential when network delays must be minimized, which has been proven by Gallager. In essence, minimum-delay routing can be achieved or approximated only by using multiple available paths to reach any one destination. In addition, using multiple paths is critical for providing fault-tolerant routing in very large networks or internetworks.

Unfortunately, there are many limitations to today's Internet routing protocols. The widely deployed routing protocol RIP [1] provides only one next-hop choice for each destination and does not prevent temporary loops from forming. Cisco's EIGRP[2] ensures loop-freedom but can guarantee only a single loop-free path to each destination at any given router. The link-state protocol OSPF[3] offers a router multiple choices for packet-forwarding only when those choices offer the minimum distance. When there is fine granularity in link costs metric, perhaps for accuracy, there is less likelihood that multiple paths with equal distance exist between each source-destination pair, which means the full connectivity of the network is still not used for load-balancing. Also, OSPF and other algorithms based on topology-broadcast (e.g., [4, 5]) incur too much communication overhead, which forces the network administrators to partition the network into areas connected by a backbone. This makes OSPF complex in terms of router configuration required.

With the exception of the diffusing algorithm for shortest multipaths (DASM), none of the routing algorithms reported before the start of this project supported multiple loop-free paths at every instant, and there existed no link-state algorithms that provided both multiple paths and loop-freedoms.

To address the limitations of today's Internet routing protocols, we developed several novel algorithms for routing of packets over multiple paths that need not be of equal cost. Formally, let a computer network be represented as a graph $G = (N,L)$, where N is set of nodes (routers) and L is the set of edges (links), and let N^i be the set of neighbors of node i . The problem consists of finding the successor set at each router i for each destination j, denoted by $S_j^i \subseteq N^i$, so that when router i receives a packet for destination j, it can forward the packet to one of the neighbor routers in the successor set S_j^i. By repeating this process at every router, the packet is expected to reach the destination. If the routing graph SG_j, a directed subgraph of G, is defines by the link set $\{(m,n) | n \in S_j^m, m \in N\}$, a packet destined for j follows a path in SG_j, be free of loops, at least when the network is stable, because routing loops degrade the network performance. In a dynamic environment, a stricter requirement is that SG_j be loop-free at *every instant*, i.e., if S_j^i and SG_j are parameterized by time t, then $SG_j(t)$ should be free of loops at any time t. If there is at most one element in each S_j^i, then SG_j is a tree and there is only one path from any node to j. On the other hand, if S_j^i's have more than one element, then SG_j is a directed acyclic graph (DAG) and has greater connectivity than a simple tree enabling traffic load balancing.

We developed fault-tolerant and self-organizing routing algorithms that provide multiple loop-free paths to each destination using only distances to destinations, the distance and second-to-last hop of the path to each destination, or partial link-state information corresponding to those links in the paths used to reach destinations.

We also introduced a generalization of loop-freedom conditions for routing algorithms based on any type of information, and applied multipath routing algorithms to a load-balancing routing framework to obtain "near-optimal" delays. A key component of this framework is a fast responsive routing protocol that determines multiple successor choices for packet forwarding, such that the routing graphs implied by the routing tables are free of loops even during network transitions. By load-balancing traffic over the multiple next-hop choices, congestion and delays are reduced significantly.

The following papers describe our results on multipath routing and minimum-delay routing:

- S.Vutukury and J.J. Garcia-Luna-Aceves, "A Distributed Algorithm for Multipath Computation," *Proc. IEEE Globecom '99*, Rio de Janeiro, Brazil, December 5-9, 1999.
- S.Vutukury and J.J. Garcia-Luna-Aceves, "An Algorithm for Multipath Computation using Distance-Vectors with Predecessor Information", *Proc. IEEE IC3N 99*, Boston, Massachusetts, October 11-13, 1999.
- S.Vutukury and J.J. Garcia-Luna-Aceves, "A Simple Approximation to Minimum-Delay Routing", *Proc. ACM SIGCOMM 99*, Cambridge, Massachusetts, September 1-3, 1999.
- J.J. Garcia-Luna-Aceves, S. Vutukury, and W.T. Zaumen, "A Practical Approach to Minimizing Delays in Internet Routing Protocols", *Proc. IEEE ICC '99*, Vancouver, Canada, Jul 6-10, 1999.

When multiple paths to destinations are provided at the routing layer, such end-to-end protocols as TCP may suffer performance degradations due to packets being delivered out of order. To solve this problem without having to establish virtual circuits in routers or tags in packets, we developed a traffic engineering approach that allows routers to forward packets of a given TCP connection over the same path, while distributing packets of different TCP flows over different paths. This work is described in the following paper:

- S.Vutukury and J.J. Garcia-Luna-Aceves, "A Traffic Engineering Approach based on Minimum-delay Routing", *Proc. IEEE IC3N 2000*, Las Vegas, Nevada, October 16-18, 2000.

It is now widely accepted that explicit resource reservations must be made in the Internet to provide the kind of guarantees (bandwidth, delay and delay-jitter) new application demand. There are two QoS architectures being proposed in the Internet today. The Integrated Services (Intserv) [6, 7] architecture provides deterministic guarantees to individual flows by reserving resources on a single route from the source to the destination using a signaling protocol (e.g., RSVP [8]); however, it cannot scale well because of the excessive state maintained in routers. The Differential Services (Diffserv) architecture [10, 9] aggregates routing and reservation state in the routers to achieve scalability, but cannot provide deterministic guarantees. Both approaches also suffer from the inherent limitation of relying on single-path routing and single-path signaling for resource reservations.

In this project, we introduced a multipath routing framework for the provision of QoS guarantees in computer networks, without the need to maintain per-flow state at routers. This is the first routing architecture capable of providing deterministic guarantees in wired networks using the same amount of state as the Diffserv architecture. This work is described in the following papers:

- S.Vutukury and J.J. Garcia-Luna-Aceves, “A Multipath Framework Architecture for Integrated Services”, *Proc. IEEE Globecom 2000*, San Francisco, California, November 27 – December 1, 2000.
- S.Vutukury and J.J. Garcia-Luna-Aceves, “A Scalable Architecture for Providing Deterministic Guarantees”, *Proc. IEEE IC3N 99*, Boston, Massachusetts, October 11-13, 1999.

3. MULTICAST ARCHITECTURES AND PROTOCOLS

In this project, we analyzed the adequacy of the existing IP multicast architecture in support large-scale applications requiring multi-point communication support. Our study shows that the current IP multicast architecture does not provide adequate support. This work is described in the following paper:

- B. Levine, J. Crowcroft, C. Diot, J.J. Garcia-Luna-Aceves, and J. Kurose, "Consideration of Receiver Interest for IP Multicast Deliver", *Proc. Infocom 2000*, Tel-Aviv, Israel, March 26-30, 2000.

There are four fundamental threats to computer communications as defined by Ford [13]. These threats are: *information leakage*, through which unauthorized receivers are able to determine information about the data or nature of the data being sent; *integrity violation,* through which messages that are being sent are altered in some manner; *denial of service*, which occurs when an attacker is able to prevent some group of legitimate users from receiving the communication service; and *illegitimate use*, which allows unauthorized users access to a service. These general threats apply to multicast data as well, though in manners different than unicast transmission. Each of these general threats can be brought about by an *enabling threat*. An enabling threat is a specific attack or occurrence that can lead to the realization of one of the four primary threats.

One well recognized enabling threat that leads to information leakage in both unicast and multicast is *eavesdropping*, where an attacker is able to listen to traffic and intercept passing data and control information. In fact, eavesdropping is probably a more serious threat for multicast than unicast; the multicast tree can extend to many more receivers and traverse many more links than a unicast path, giving the attacker more places to listen to the multicast traffic. This property of the multicast tree, that is possibly extends over many more links in the network then a single unicast path, makes other enabling threats more dangerous as well. The possibility of *traffic analysis*, through which an attacker gains some useful information by analyzing which network members send or receive which information, is also increased since there are more possible points on the network to gather information. There are also more points for an attacker to intercept and alter information, leading directly to the possibility of integrity violation. The effect of altering data or control packets in a multicast group can be compounded by the very nature of the service; a single altered packet can be copied repeatedly on its way across the tree, arriving at many different receivers and in essence amplifying the original attack. Besides allowing rapid spread of altered information, the fact that messages are copied as they traverse the tree allow for very effective denial of service attacks. By injecting large numbers of spurious packets onto the tree a single attacker could possibly overwhelm many receivers; this is a strong argument for doing data packet verification at the routing level rather than the application level. In addition, in a shared multicast tree an attacker could induce loops in the structure of the tree by altering or replaying control messages; loops have the effect of saturating the links they traverse, essentially denying service across those links.

All of the above threats and attacks are most easily realizable if the attacker is able to gain illegitimate access to the network. This can occur in different ways. An attacker may be able to *masquerade* as an authorized router. In this case they will be privy to all data communications and be able to issue control packets to the group with the same authority as the subject of the masquerade. An attacker may also use a *man in the middle* attack by masquerading as two nodes at once, each communicating with the other. By intervening in the communication between the two nodes it is impersonating, the attacker can gain any information that might pass between them, including any shared keys. A masquerade may also be perpetuated by altering information in or coming from a trusted key server. In a multicast tree, it is difficult to detect and prevent illegitimate use. The anonymous nature of multicast, in which one sender can communicate with many receivers without knowing who or where they are, make it difficult to verify that the receivers are authorized. Whereas for a unicast transmission there are only two participants, one at either end of the path and a sender can verify the identity of the single receiver easily, in multicast there can be a large number of participants and authentication is made more difficult. A single sender could be overwhelmed if it were required to authenticate many receivers. In multicast, the membership of the group can also change at any time, so whatever mechanism exists for performing authentication must be able to handle member joins and deletions effectively, without disturbing the rest of the group. This calls for a distributed and efficient authentication system.

The approach used in designing a security framework for a multicast routing protocol depends on what assumptions are made about the nature of the threat, on the nature of the network environment and on the nature of the multicast protocol being used. It is only prudent to assume that the threat includes all those described above; that an attacker can access the network at any given point and inject, alter or replay control or data traffic or that an attacker will attempt to become part of the multicast group by impersonating whatever entity makes entry into the group possible. The nature of the network also plays a large role in the design of the multicast routing security protocol. If the network were composed of entirely trusted routers, it would be possible to verify all paths across the tree on a hop-by-hop basis. If the routers were not all trusted, then some mechanism for verifying a receiver across an untrusted path must be included. Since many multicast protocols rely on the existing unicast routing, then if an attacker were able to alter the unicast routing tables they could force a branch of the multicast tree to pass through any compromised router they desired. The security protocol must be able to deal effectively with all these conditions.

In this project, we developed Keyed HIP (KHIP), a secure multicast routing protocol based on hierarchical multicast routing. The multicast routing protocol we used as the basis of KHIP is a hierarchical multicast routing protocol called HIP. HIP allows the used of shared tree protocols such as PIC-SM or CBT as the inter-domain routing protocol in a hierarchy that can include any routing protocol at the lowest level. The architecture consists of two protocols; one that encapsulates an entire routing domain to allow it to appear as the *virtual router* on a higher-level shared tree; and a second protocol that provides mechanisms for rendezvous point or core distribution and recursively applies the first protocol to produce trees of domains that contain trees of domains. HIP is the first architecture to allow any multicast protocol at the lowest level while using a shared tree for higher-level routing. It provides a simple, efficient mechanism for RP or core location dissemination. HIP aligns easily with existing unicast domains and it does not require explicit assignment of levels except at the highest level. HIP is suitable for any shared tree protocol, such as PIM-SM or CBT, that forms a tree by sending join messages to some central router. It can also provide additional robustness for the shared tree protocol through the ability to replace a single rendezvous point or core with several routers that operate together in a distributed fashion and can tolerate members failing.

The advantage of using KHIP for secure multicasting is that the hierarchical structure allows authentication and verification procedures to take place within and be limited to one domain. Using public-key cryptosystems is fairly simple and effective if the number of keys that needs to be stored is small; a master router within each domain could easily track the number of keys required to authenticate those within his domain. As receivers pass from domain to domain, key information could be shared within the higher-level domain so that an extensive global system of obtaining and verifying public keys was not necessary. The authentication message could also carry a shared key for a faster, symmetric data encryption mechanism for use only within that domain. Re-keying a domain is much easier than re-keying the entire multicast group, and as each sub-domain would use its own shared key, the re-keying is limited to a single level. Membership of a domain is also easier to track; the master router can exchange "heartbeat" messages with each receiver in the domain to make sure that members leaving the group cause a re-keying of the domain. Data flow into or out of the domain can be checked to make sure that the sender is authorized and that the data within the packet has not been altered. This will limit the scope of a denial of service attack to at most a single domain. In addition, different procedures can be followed within each domain based on the nature of the network within the domain; some domains might contain all trusted routers and run a secure routing protocol and allow verification of the entire path, while others might run an insecure routing protocol and require stricter methods of verification across untrusted routers. The hierarchical nature allows for a variety of security levels and network conditions.

We showed that other shared-tree multicast routing protocols are subject to attacks against the multicast routing infrastructure that can isolate receivers or domains or introduce loops into the structure of the multicast routing tree. KHIP changes the multicast routing model so that only trusted members are able to join the multicast tree. This protects the multicast routing against attacks that could form branches to unauthorized receivers, prevents replay attacks and limits the effects of flooding attacks. Untrusted routers that are present on the path between trusted routers cannot change the routing and can mount no denial-of-service attack stronger than simply dropping control messages. KHIP also provides a simple mechanism for distributing data encryption keys while adding little overhead to the protocol.

The work on KHIP is described in the following paper:

- C.Shields and J.J. Garcia-Luna-Aceves, "A Scalable Protocol for Secure Multicast Routing", *Proc. ACM SIGCIMM 99*, Cambridge, Massachusetts, September 1-3, 1999.

4. RELIABLE BROADCASTING

Today's network architectures are based on either passive datagram handling in which each packet is handled the same way and carries the same set of instructions for the "static code" running in the routers, and passive connection-oriented handling in which a connection is established to determine the way in which subsequent packets for the same connection will be handled. All routing and multicasting protocols require the *reliable dissemination* of routing-table information in one way or another, and this information reflects the same physical network topology; protocols like OSPF based on topology broadcast are the simplest to understand, because the same topology map is disseminated to each node. In contrast, with active packets, a node can be instructed to treat packets differently, depending on many factors, including user groups and the location of the handling node. Furthermore, it is also possible to establish virtual topologies on top of the physical topology to expedite packet forwarding; this is a generalization of establishing spanning trees in bridge-based internets or establishing shared routing trees for multicasting in an internet.

We investigated protocols needed for the control of active networks that build virtual topologies to achieve different router behavior over different such topologies. We started the development of an *internetwork reliable concast protocol* (IRCP), whose objective is to support the trusted distribution of active packets to allow routers to build virtual topologies. The two types of information exchange protected by the protocol are: (a) active packets communicated between neighboring nodes and (b) active packets multicast among an arbitrary set of routers.

Together with SRI International, we verified a new reliable broadcast protocol to serve as the basis for IRCP. All reliable broadcast protocols based on flooding that have been proposed for dynamic topologies in the past (e.g., [18]) are based on the routing protocol by Merlin and Segall [17]. These protocols proceed in cycles triggered and terminating at the source of the message. A cycle consists of two phases. First, the message propagates one additional hop away from the source, then acknowledgments to the message propagate towards the source. The source starts by sending a message that is acknowledged by all its neighbors. When the source receives the acknowledgments from its neighbors, it then resends the message asking the neighbors to propagate the message one more hop to their own neighbors. The neighbors forward the message to their neighbors and send the acknowledgments back to the source when they receive the acknowledgments from all their own neighbors, and so forth. This scheme incurs too much communication overhead to be attractive for a wireless network; furthermore, an implicit assumption of this approach is that a node can have a fairly stable successor to the source, which does not apply in a network with mobile nodes.

The reliable broadcast protocol (RBP) addressed in this project ensures that every node connected to the source node receives the information from the source node at least once, and that the source node is positively informed that the information reaches all the connected nodes in the network within a finite time. RBP works in a similar way to PIF for the case of a static network. However, in contrast to prior approaches of reliable broadcasting in dynamic networks, or proposed protocol requires the source to send each broadcast message only once, and diffusing computations [16] are used to eliminate the need for the source node to control the propagation of information in multiple rounds. Instead of defining a single successor for each node in a directed acyclic graph (DAG), RGP permits each node to define a successor set formed with all those neighbors who transmit the source's message. The work on the verification of RBP is described in the following paper:

- G. Denker, J.J. Garcia-Luna-Aceves, et al, "Specification and Analysis of a Reliable Broadcasting Protocol in Maude", *Proc. 37th Allerton Conference on Communications, Control, and Computing*, September 22-24, 1999.

5. CONCLUSIONS

This project made a number of contributions to advance the state of the state of the art in internet-working. Key contributions of this project include new solutions for routing over multiple paths, secure multicasting, destination-based provisioning of quality of service, and reliable broadcast of packets that can be used to down-load code in routers.

The progress made in multipath routing enables new research on fault-tolerant routing in very large networks, internetworks and sensor networks. Of particular interest is the provision of fault-tolerant routing with multiple constraints resulting from the environment or the desired use of information. In general, a new architecture and protocols for fault-tolerant internetworking can be developed such that: (a) routers can *protect* efficiently against attacks and faults, and *detect and respond* to them in a timely manner; (b) no routing and multicasting function has single point of failure; and (c) QoS guarantees are provided in a scalable and fault-tolerant manner. Our results on secure multicasting were the first to address securing the routing infrastructure itself, and serve as a benchmark for future work. Our work on reliably broadcasting of control packets is only a first step in the development of an architecture for the distribution of code or instructions to routers in a trusted and fault-tolerant manner.

6. REFERENCES

[1] C. Hendrick. Routing Information Protocol. RFC, 1058, June 1988.

[2] R. Albrightson, J.J. Garcia-Luna-Aceves, and J. Boyle. EIGRP-A Fast Routing Protocol Based on Distance Vectors. *Proc. Networld/Interop 94*, May 1994.

[3] J. Moy. OSPF Version 2. RFC, 1247, August 1991.

[4] J. Spinelli and R. Gallager. Event Driven Topology Broadcast without Sequence Numbers. *IEEE Trans. Commun*., 37:468-474, 1989.

[5] R. Perlman. Fault-tolerant broadcast of routing information. *Computer Networks and ISDN,* 7, 1983.

[6] S. Shenker, D. Clark, and L. Zhang. A service model for an integrated services internet. *Internet Draft,* Oct. 1`993.

[7] E. Crawley et al. A Framework for qos-based routing in the internet. *Internet Draft*, April 1998.

[8] L. Zhang et al. RSVP: A New Resource Reservation Protocol. I*EEE Communications Magazine*, 31 (9): 8-18, 1993.

[9] Y. Bernet et al. An Framework for Differentiated Services. *Internet Draft,* May 1998.

[10] D. Black et al. An architecture for Differentiated Services. *Internet Draft*, May 1998.

[11] T. Ballardie et al., "Core Based Trees (CBT) – An Architecture for Scalable Inter-Domain Multicast Routing," *Proc. ACM SIGCOMM 93,* San Francisco, CA, 1993.

[12] S. Deering, et al., "An Architecture for Wide-Area Multicast Routing," *Proc. ACM SIGCOMM 94*, London, UK, August 1994.

[13] W.Ford. *Computer Communications Security*, Prentice Hall, 1994.

[14] D.D. Clark, S. Shenker, and L. Zhang, "Supporting Real-Time Applications in an Integrated Services Packet network: Architecture and Mechanisms," *Proc. ACM SIGCOMM 92*, August 1992.

[15] D. Ferrari, A. Banerjea, and H. Zhang, "Network Support for Multimedia – A Discussion of the Tenet Approach," *Computer Networks and ISDN Syst*. Vol. 26, No. 10, pp 1167-1180, 1994.

[16] Edsger W. Dijkstra and C.S. Scholten. Termination detection for diffusing computations. *Information Processing Letters*, 11 (1):1-4, June 1980.

[17] P.M. Merlin and A. Segall, "A Failsafe Distributed Routing Protocol", *IEEE Trans. Commun.*, Vol. 27, No. 9, pp 1280-1288, 1979.

[18] A Segall and B. Awerbuch, "A Reliable Broadcast Protocol", *IEEE Trans. Commun*., Vol. 31, No. 7, ppd. 896-901, 1983.

[19] A. Segall, "Distributed Network Protocols", *IEEE Trans. Info. Theory*, Vol.29, No.1, pp.25-35, 1983.l